

IN THE SPECIFICATION:

*Please amend the paragraph at page 9 lines 11-25 as follows:*

When there are more jobs to run than high-performance cores available in a heterogeneous system with dynamic job assignment, an operating system associated with the processor core pool dispatches the execution of application programs to various processor cores and monitors their resulting performance. Alternatively, such control is implemented in firmware or special-purpose hardware. Such performance can be measured with simple metrics such as instructions per second. Metrics are collected on how well an application runs on a particular processor core in the pool, for example during a one millisecond period. After several milliseconds of swapping execution of jobs between cores and monitoring their resulting performance, the operating system can build a table with the relative performance of each job on each type of core.

*Please amend the paragraph at page 15 line 32 through page 16 line 12 as follows:*

In general, a target processor core from a pool being sought for a job is the one that yields acceptable performance. A balance with other measures can be determined statically for each workload based on data from prior executions of the workload. Alternatively, such balance may be obtained dynamically by empirically determining it at run-time. Metrics are collected on how well an application runs on a particular processor core in the pool, for example during a one millisecond test period. If the current processor core is yielding better results than a previous processor core, then the job will not be transferred, and will be allowed to continue executing. If not, the job can be returned to the previous processor core in the ordered pool or a next processor

core can be tried. The resource requirements between application programs can vary, as well as the requirements at different times within a single application.